Software-Defined Radio: Learned Trust-Weighted Collaborative Navigation Under Satellite Jamming and Spoofing

Gaurav Sharma[1,2]*, C.R.S. Kumar[1]†, R. Rajesh[2]‡

[1] Defence Institute of Advanced Technology (DIAT), Girinagar, Pune, Maharashtra 411025, India

[2] Centre for Airborne Systems (CABS), Defence Research and Development Organisation, Wind Tunnel Road, Bengaluru, Karnataka 560064, India

** Doctoral Research Scholar, Department of Computer Science and Engineering, Defence Institute of Advanced Technology, Girinagar, Pune, Maharashtra, India*

† Professor, Department of Computer Science and Engineering, Defence Institute of Advanced Technology, Girinagar, Pune, Maharashtra, India

‡ Scientist G, Centre for Airborne Systems, DRDO, Wind Tunnel Road, Bengaluru, India

Abstract

Global Positioning System signals used by airborne platforms for position, navigation, and timing are vulnerable to jamming and spoofing, motivating cooperative navigation architectures that do not depend on satellite signals alone. This paper presents a simulation-validated framework for collaborative navigation over a tactical software-defined radio network, encompassing automated detection of satellite signal denial, secure discovery of cooperating platforms, network-wide time synchronization, several independent methods for determining own-ship position from cooperating aircraft and ground sites, and automatic recovery once satellite signals are restored. The central contribution is a learned, attention-based fusion mechanism that assigns trust to each navigation source according to context, extending an architecture previously developed for

airborne radar track fusion. The network was trained on many simulated denial events spanning several scenario geometries and both jamming and spoofing conditions, and was validated across multiple independent training runs. Relative to a classical weighted-least-squares baseline, the trained network reduced median position error substantially, with the largest improvement in the most geometrically challenging conditions. A sequential filter baseline representative of standard navigation-literature practice was also implemented and found considerably more vulnerable to spoofing than either the classical or learned alternative, attributed to its lack of an independently referenced integrity check.

Keywords: collaborative navigation, GPS-denied navigation, GPS spoofing, sensor fusion, attention mechanism, trust-weighted fusion, software-defined radio, cooperative PNT, Extended Kalman Filter, integrity monitoring.

Nomenclature

d = position error

N = number of independent training runs

$PDOP$ = Position Dilution Of Precision

σ = standard deviation of a measurement or estimate

w = trust weight assigned to a navigation source

I. INTRODUCTION

Our Modern airborne operations depend heavily on satellite-based Position, Navigation, and Timing (PNT) services, which underpin mission-critical functions including surveillance, sensor fusion, weapon guidance, and network-centric command and control. Recent conflicts have demonstrated that the electromagnetic spectrum is now a contested domain where GPS jamming and spoofing are widely and effectively used to degrade or deny satellite navigation to airborne platforms, resulting in a single-point-of-failure risk that motivates cooperative, multi-source alternatives that do not rely on GPS alone [1–6].

A representative operational concept for addressing this problem is as follows. Under normal conditions, each platform navigates using its own onboard GPS and inertial navigation system (INS). On detection of degradation, denial, or spoofing of its GPS signal, a platform automatically initiates a Collaborative PNT service over a tactical software-defined radio (SDR) network. On the other hand, the Platforms that continue to possess a reliable, authenticated GPS solution periodically broadcast secure PNT information together with quality metrics. A GPS-denied platform validates the received information, assesses the quality of each available navigation source, and integrates the external information with its own INS solution through appropriate estimation and fusion algorithms. Independently derived position fixes obtained via inter-platform ranging and bearing, relayed SDR radar-track correlation, and passive Radar Warning Receiver / Electronic Support Measures (RWR/ESM) geolocation are used both to improve fusion accuracy and to cross-check the integrity of the received cooperative data. Upon restoration of onboard GPS, the platform automatically validates the recovered signal and transitions back to onboard navigation without disturbing ongoing operations.

This paper establishes a full-stack simulation framework executing the concept of operations, together with a learned fusion mechanism that is the first application of a variable-cardinality attention architecture originally developed for radar/IFF track association and adaptive covariance-intersection fusion in our prior work [7,8] to the navigation-source fusion problem. Unlike a constant fusion head, the attention mechanism naturally accommodates a dynamic cooperative source set. The size and composition of this head change from step to step as platforms enter and leave range, as ranging anchors become available or unavailable, and as denial conditions evolve.

A. Contributions

(1) A simulation framework covering the full functional scope of Collaborative PNT over a tactical SDR network comprising of GPS health monitoring, cooperative platform discovery, network time synchronization, navigation information distribution, own-

position determination via ranging/multilateration/bearing/relayed radar tracks/passive RWR-ESM ranging, navigation source management, fusion, continuous integrity monitoring, and automatic recovery.

(2) A three-layer relayed-track correlation algorithm for identifying own-ship among third-party radar tracks: direct SDR network-identifier matching, IFF-code matching, and INS-gated Mahalanobis association as successive fallbacks, validated at 100% correlation success against as many as eighteen simultaneous distractor aircraft.

(3) A passive geolocation capability for RWR/ESM that adds RSSI-based ranging to bearing-only tracking for known, friendly emitters, with a validated bias-correction for the log-normal estimation bias inherent to naive received-power inversion.

(4) An attention-based, variable-cardinality trust-weighted fusion network, trained on a large denial-step dataset and validated across five independent training runs, that improves median position error over a classical weighted least-squares (WLS) baseline within the range of 7 to 57% depending on scenario geometry and denial mode.

(5) A literature-standard Extended Kalman Filter (EKF) baseline, which discloses that sequential filtering without an independently referenced integrity check is substantially more vulnerable to spoofing than either a memoryless per-step solution or the proposed learned fusion including a negative ablation result showing that naive chi-square innovation gating does not resolve this vulnerability.

The remainder of this paper is organized as follows. Section II summarizes the concept of operations and functional requirements underlying the framework. Section III reviews related work in cooperative navigation, classical and learned multi-sensor fusion, and GNSS integrity monitoring. Section IV describes the system architecture. Section V describes the experimental methodology, including dataset generation and the multi-seed training and validation protocol. Section VI presents results. Section VII discusses their

interpretation and implications. Section VIII states limitations and future work, and Section IX concludes.

II. CONCEPT OF OPERATIONS AND FUNCTIONAL REQUIREMENTS

The framework implements the functional requirements summarized in Table 1, organized around nine functional groups: (1) GPS health monitoring, including jamming and spoofing detection and position-dilution-of-precision (PDOP) based fix rejection; (2) cooperative platform discovery and authentication; (3) network-wide time synchronization; (4) cooperative navigation information distribution, modeled as an authenticated, quality-tagged, genuinely multicast broadcast delivered after a bounded, real dissemination delay; (5) own-position determination via SDR ranging/Multilateration, bearing-only triangulation, relayed radar/IFF track correlation, and passive RWR/ESM ranging and bearing; (6) navigation source quality assessment and trust-based source management; (7) navigation fusion, both classical and learned; (8) continuous integrity monitoring; and (9) automatic, bump-less recovery upon GPS reacquisition.

Two functional areas are deliberately modeled at the level of their received output rather than their underlying protocol mechanics, consistent with the scope of a navigation-algorithm research framework rather than an RF/cryptographic engineering effort. First, the “secure” property of cooperative PNT broadcasts is represented as an authenticated, quality-tagged message structure delivered to every in-range receiver; the cryptographic and RF mechanisms that would produce this property in a fielded system are out of scope. Second, IFF identification is represented via a numeric code exchange rather than a full Mode-5 cryptographic protocol. Both scoping decisions are stated explicitly here and in Section VIII to avoid overclaiming security properties that were not implemented or tested.

Functional Area	**Representative Capability**
GPS Health Monitoring	Jamming/spoofing detection; PDOP-based rejection gate; clock-instability flagging
Cooperative Platform Discovery	Authenticated peer discovery; MANET multi-hop connectivity (BFS-based)
Time Synchronization	Periodic network-wide correction bounding long-run clock-bias

	drift
Navigation Info. Distribution	Multicast, quality-tagged, latency-delayed PNT broadcast
Own-Position Determination	SDR ranging/Multilateration; bearing triangulation; 3-layer radar/IFF track correlation; passive RSSI-based RWR/ESM ranging
Source Management	Per-source quality scoring (accuracy, latency, confidence, authentication)
Navigation Fusion	Classical WLS baseline; attention-based learned trust-weighted fusion (this work)
Integrity Monitoring	Chi-square consistency gating with persistence-based isolation/reinstatement
Automatic Recovery	Validated, bump-less transition back to onboard GPS on reacquisition

***Table 1**. Functional scope of the implemented framework.*

III. RELATED WORK

A. Cooperative and Relative Navigation

Cooperative or relative navigation improves individual navigation accuracy by sharing information among platforms, and has been demonstrated across ground, aerial, and space domains using ranging, bearing, and visual measurements [1–6]. Ellingson et al. [4] introduced a relative-navigation framework for single fixed-wing GPS-denied flight and later extended it to a decentralised, cooperative multi-aircraft setting with a coordinated-reset mechanism [5], demonstrating navigational improvement from cooperation on small fixed-wing unmanned aircraft directly relevant prior work published in this journal. Shen et al. [6] similarly showed that peer-to-peer ranging can meaningfully augment degraded GNSS in multi-UAV formations. The present work advances this body of literature in two important respects. First, rather than focusing on a single measurement modality or a narrowly defined cooperative-navigation configuration, it integrates a substantially richer and heterogeneous set of cooperative source types within one framework. Second, it replaces fixed, model-dependent fusion weights with a learned, context-adaptive attention mechanism that can adjust source contributions as platform availability, source quality, network topology, and denial conditions change. This combination directly addresses the variable and nonstationary conditions that limit conventional fixed-gain fusion in contested environments.

B. Classical and Learned Multi-Sensor Fusion

Traditional multi-sensor fusion commonly relies on Extended Kalman Filters and factor graphs, which use static or linearized observation models that may perform poorly under the non-Gaussian noise typical of contested electromagnetic environments. Covariance Intersection (CI) [9] is a long-established, conservative approach to fusing estimates with unknown correlation, and underlies the classical fusion baseline used in this work. Sequential (Kalman-filter-based) and federated architectures remain the dominant classical approach to multi-sensor integrated navigation [10,11], while factor-graph optimization has emerged as a competitive alternative, particularly where sensors are asynchronous or intermittently available [11,12]. Deep-learning-based fusion, and attention-based fusion specifically, has recently been applied to inertial-sensor and multi-modal navigation problems [13,14], generally to combine fixed, small sets of co-located sensors on a single platform. In our own prior work on airborne radar/IFF track fusion, a transformer-based attention mechanism was developed for associating asynchronous radar and IFF detections to system tracks [7], and subsequently extended into an end-to-end framework in which the fusion stage itself learns per-scan trust weights via adaptive Covariance Intersection rather than relying on fixed mixing coefficients [8]. This work applies attention based fusion to a variable-cardinality, multi-platform source set whose size and composition changes at each timestep with changing network topology and denial conditions. It adapts an architecture originally developed for track level sensor fusion rather than designing a fusion network from first principles.

C. GNSS Integrity Monitoring and Spoofing Detection

Receiver Autonomous Integrity Monitoring (RAIM) and its INS-coupled extensions remain the standard approach for detection of GNSS faults and spoofing via consistency checking of redundant measurements [15]. A repeated theme in this literature, echoed in the present work's own findings (Section VI.C), is that integrity monitoring is only as robust as the independence of its reference. An INS-referenced check can itself be compromised by a

slow-ramp spoofing attack that the INS has already partially absorbed before detection of spoofing, motivating research , that has resulted us presenting this framework , that cross-check against an independently and freshly computed cooperative solution rather than the platform's own potentially-compromised state estimate is a good approach to operate in GPS jamming/spoofing conditions.

IV. SYSTEM ARCHITECTURE

A. Overview

The framework is organized as a sequential per-timestep pipeline emulating the functionalities as defined in Table 1: GPS sampling and health evaluation, cooperative broadcast publication and reception (as per a modeled dissemination latency), periodic network time synchronization, own-position determination via SDR ranging, bearing, relayed radar/IFF track correlation, and passive RWR/ESM ranging, navigation source scoring and filtering, fusion (classical or learned); continuous integrity monitoring; and automatic recovery. Figure 1 presents this pipeline as implemented. Five scenario layouts were implemented by varying anchor geometry and platform density i.e a compact loose formation, a widely distributed anchor layout, a deliberately poor (clustered/near-collinear) anchor geometry, a large-area layout exceeding 250 km spanning a two-hop mobile ad hoc network (MANET) relay, and a 20-aircraft/10-ground-site “large fleet” layout with randomized platform placement and randomized jamming-target selection.

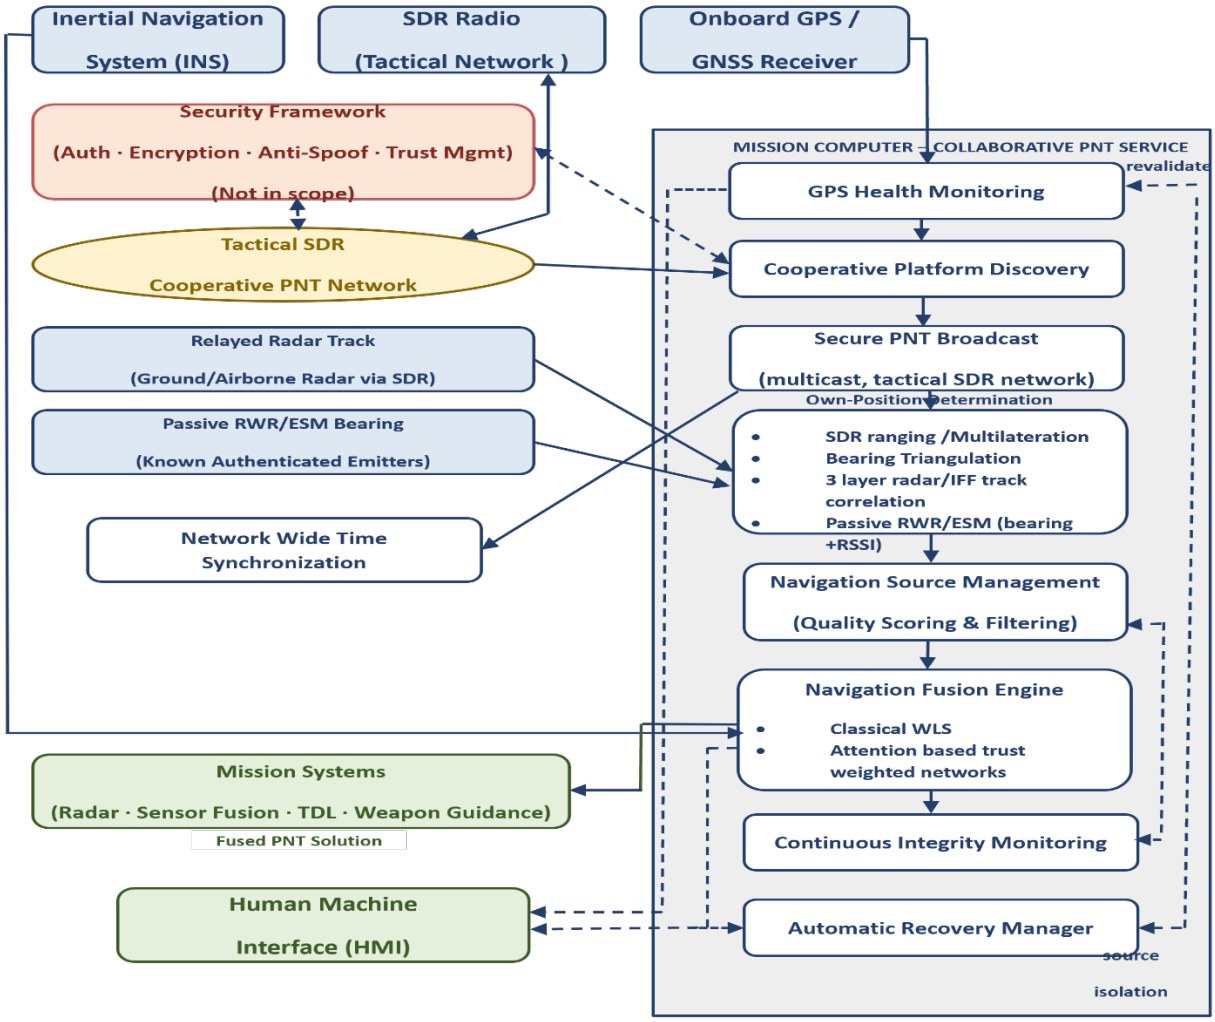

Fig. 1*. System architecture of the implemented Collaborative PNT framework, organized around the nine functional groups of Table 1.*

B. Relayed Radar-Track Correlation

Determining which of several system tracks corresponds to one's own aircraft is addressed with three successive, independently-probabilistic mechanisms. First, in the SDR environment all SDR enabled platforms have unique SDR track ID, which is tagged after the association and fusion is performed at the fusion centre and is distributed on to the network. By cross-referencing its own unique track ID against SDR broadcasts it independently receives that identifier is matched directly against the querying platform's

own identifier. Second, if IFF interrogation succeeds for the track, its numeric code is matched against the platform's own registered code. Third, if both the previously mentioned steps fail, an INS-gated Mahalanobis association is attempted against the platform's own predicted state; this classical fallback is structurally the same association problem matching a predicted own-ship state against a set of candidate radar/IFF detections addressed with a learned, transformer-based attention mechanism in our prior track-association work [7], and is a natural candidate for a similar learned upgrade in future work. This layered design was validated at 100% correlation success across scenarios with up to eighteen simultaneous distractor aircraft, with SDR-identifier matching resolving the substantial majority of cases and the remaining mechanisms providing graceful fallback.

C. Passive RWR/ESM Geolocation

Radio Warning Receiver / Electronic Support Measures (RWR/ESM) geolocation is extended beyond bearing-only tracking to include range estimation via received signal strength (RSSI), exploiting the fact that a known, authenticated friendly emitter's nominal transmit power, signature is registered and therefore known to the receiver an assumption that does not hold for unknown or hostile emitters, which is precisely why this technique is specific to authenticated friendly emitters. Because range enters the free-space path-loss equation logarithmically, naively inverting a noisy received-power measurement is a biased estimator (a consequence of Jensen's inequality). This bias was identified and corrected with a standard log-normal mean-unbiasing correction, verified to reduce mean estimation error to under 1% while preserving a realistic ≈47% relative range uncertainty consistent with published RSSI-ranging performance. Range and bearing to a common friendly emitter are combined into a direct polar position fix; range-only measurements to three or more emitters are combined via multilateration.

D. Attention-Based Trust-Weighted Fusion

The learned fusion network extends the source-encoder / multi-head self-attention architecture and adaptive Covariance-Intersection fusion stage previously developed for radar/IFF track fusion [8] from a fixed set of track-level sensors to a variable-cardinality set of navigation sources (inertial-only prediction, SDR multilateration, bearing triangulation, and the combined relayed-track/RWR-ESM own-position-determination fix). Each candidate source is encoded from its position (relative to the current inertial estimate) and its covariance diagonal is tagged with a learned embedding for its source type; multi-head self-attention [16] over the resulting token set, followed by a softmax normalization anchored against the inertial source, produces a set of trust weights. These weights combine the sources information matrices in a Covariance-Intersection-style information-form fusion, so that the mechanism generalizes classical inverse-covariance weighting rather than replacing it outright . The network can, in principle, recover the classical solution as a special case, and departs from it only where the training data indicates that it should.

V. EXPERIMENTAL METHODOLOGY

A. Dataset Generation

Training data were generated by running the classical fusion pipeline across all five scenario layouts under both jamming and spoofing conditions. At each denial timestep, the simulator recorded the complete set of candidate sources available to the learned network, together with the ground-truth position. Using the classical rather than learned mode kept the simulated trajectories stable and realistic while preserving the source inputs used for training. The final dataset contained 182,000 denial-timestep records: fifteen seeds for each of four scenario layouts and five seeds for the large-fleet layout, with both jamming and spoofing conditions represented.

B. Training Protocol

The network was trained with a combined loss comprising squared position error and a log-space innovation-consistency (NEES) term penalizing miscalibrated fused covariance, following the same calibration-aware training pattern used in our prior track-fusion work [7,8]. Training and validation sets were split by (scenario, seed, denial-mode) group never by individual row and stratified so that every scenario/denial-mode combination is guaranteed representation in both splits. A previous, non-stratified split was found to leave one combination entirely without held-out validation data purely by chance, which the stratified split corrects. To assess statistical sturdiness rather than relying on a single trained model, five independent training runs were performed with different (data-split, network-initialization) seed pairs, and results are reported as the mean and standard deviation across these five runs.

C. Baselines

Two classical baselines were implemented for comparison. The first, denoted Classical WLS, is the memoryless weighted-least-squares multilateration solution already used to drive the simulator's validated classical fusion mode. It recomputes a position fix from scratch at every timestep from whichever ranging anchors are currently available, with no temporal state. The second baseline is a conventional constant-velocity Extended Kalman Filter, implemented as a six-state filter with a piecewise-white-acceleration process model and sequential updates from every available cooperative source at each timestep. This baseline isolates the value of temporal state estimation while deliberately excluding the learned trust allocation and integrity-monitoring mechanisms, thereby providing a direct test of whether temporal smoothing alone can address the denial and spoofing conditions. An ablation variant adds chi-square normalized-innovation-squared (NIS) gating with persistence-based isolation and reinstatement, structurally identical to the framework's own continuous integrity monitor, to test whether a minimal integrity-monitoring addition closes the gap between the plain EKF and the learned network under spoofing.

VI. RESULTS

A. Multi-Seed Validated Comparison

Table 2 reports median position error for the classical WLS baseline, the plain EKF baseline, and the trained trust-weighted fusion network (mean ± standard deviation across five independent training runs), across all ten scenario/denial-mode combinations.

Scenario / Denial Mode	**Classical WLS (mean ± std, m)**	**EKF (plain) (m)**	**Trust-Net (ours) (mean ± std, m)**	**Improvement vs. Classical**
Distributed / JAMMED	52.2 ± 0.0	48.2	23.9 ± 0.6	+54.3%
Distributed / SPOOFED	50.7 ± 0.0	58.7	24.1 ± 1.2	+52.4%
Formation / JAMMED	25.4 ± 0.0	26.0	15.8 ± 2.0	+37.8%
Formation / SPOOFED	70.6 ± 0.0	230.8	65.4 ± 4.7	+7.4%
Large Area / JAMMED	22.9 ± 0.0	27.8	14.9 ± 2.0	+34.9%
Large Area / SPOOFED	22.4 ± 0.0	27.6	14.9 ± 1.5	+33.5%
Large Fleet / JAMMED	44.5 ± 5.3	30.3	32.4 ± 3.0	+27.2%
Large Fleet / SPOOFED	41.5 ± 0.8	29.1	31.6 ± 2.1	+23.8%
Poor Geometry / JAMMED	53.2 ± 0.0	83.7	22.8 ± 0.4	+57.1%
Poor Geometry / SPOOFED	53.4 ± 0.0	90.3	24.6 ± 0.4	+53.9%

Table 2. *Median position error by scenario and denial mode. Classical WLS values have zero reported standard deviation for the four scenarios with fifteen total seeds, since a three-seed held-out fold is large enough (4,200 records) that the median is stable regardless of which specific seeds are drawn; large_fleet (five total seeds, one held out per run) shows genuine seed-to-seed geometry variation in the classical baseline itself.*

The learned network improved median position error over the classical WLS baseline in nine of ten scenario/denial-mode combinations in all five independent training runs, and in four of five runs for the tenth (formation/SPOOFED). Figure 2 presents the full three-way comparison (classical WLS, plain EKF, and the trained network) from a representative single run. Figure 3 ranks the ten combinations by percentage improvement. Figure 4 presents the multi-seed mean ± standard deviation comparison . Figure 5 shows a representative training convergence curve.

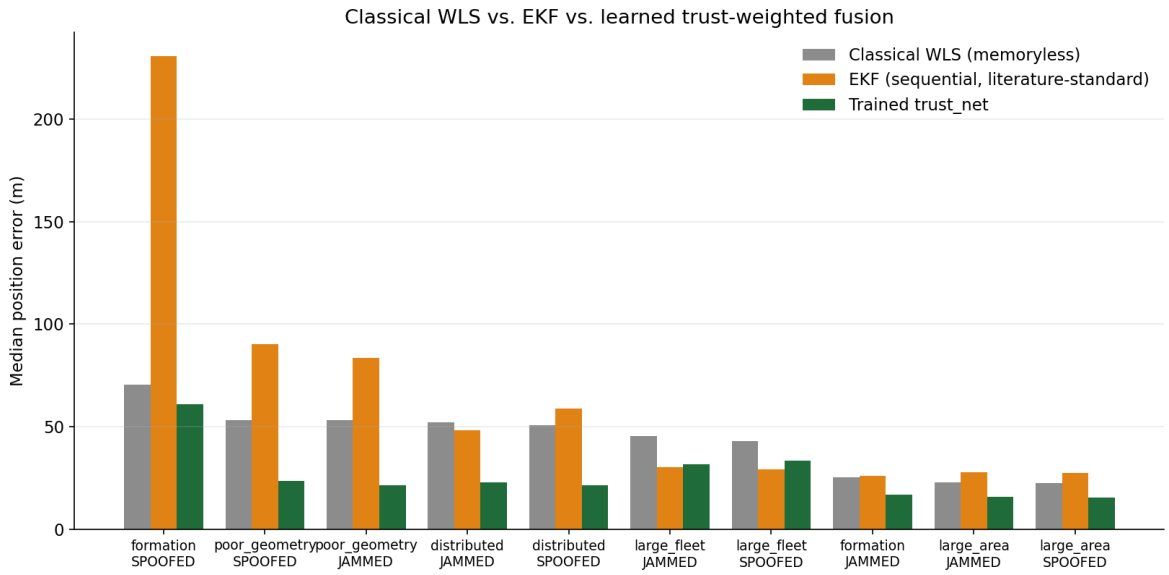

Fig. 2. *Classical WLS vs. EKF vs. trust-weighted fusion, median position error by scenario/denial mode (single representative run).*

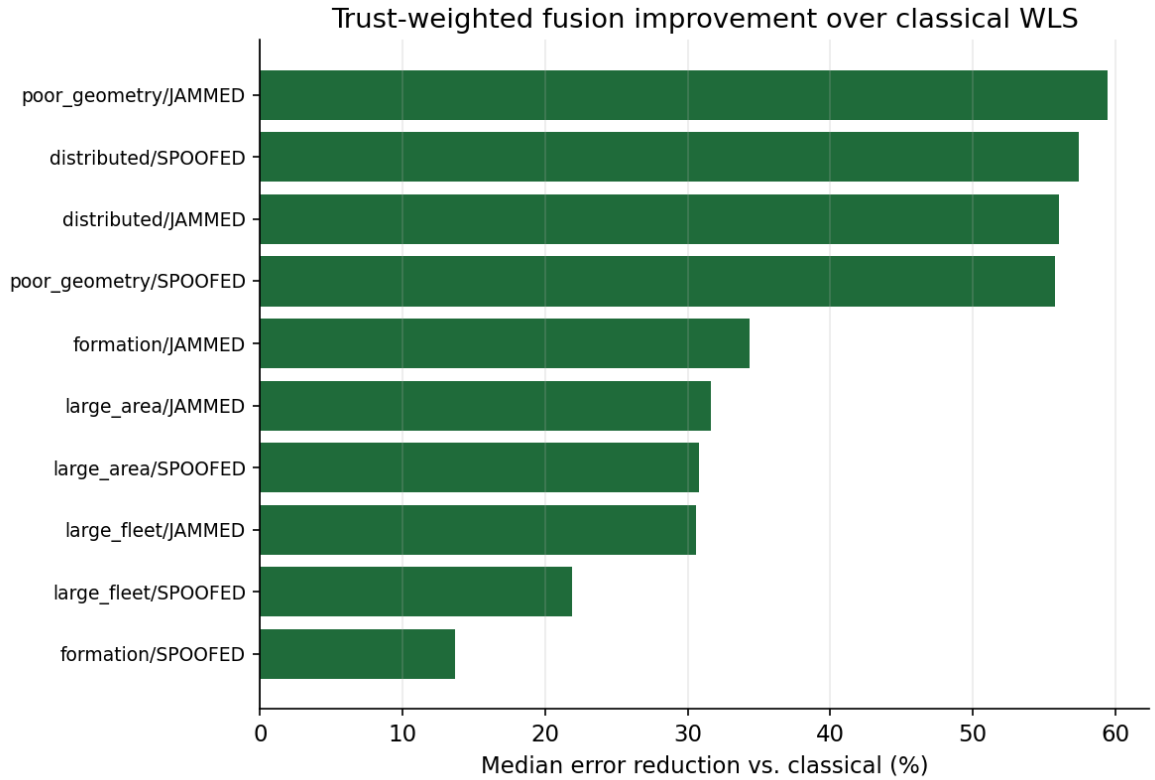

Fig. 3. *Trust-weighted fusion's percentage improvement in median position error over classical WLS, ranked by magnitude.*

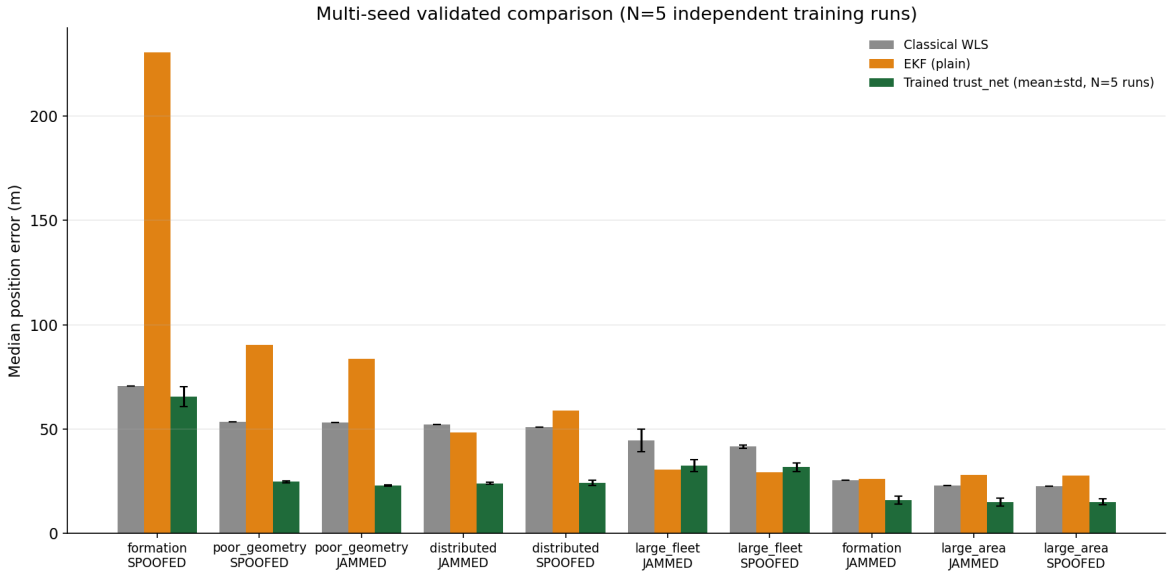

Fig. 4*. Multi-seed validated comparison (N = 5 independent training runs), classical WLS and trust-net reported as mean ± standard deviation.*

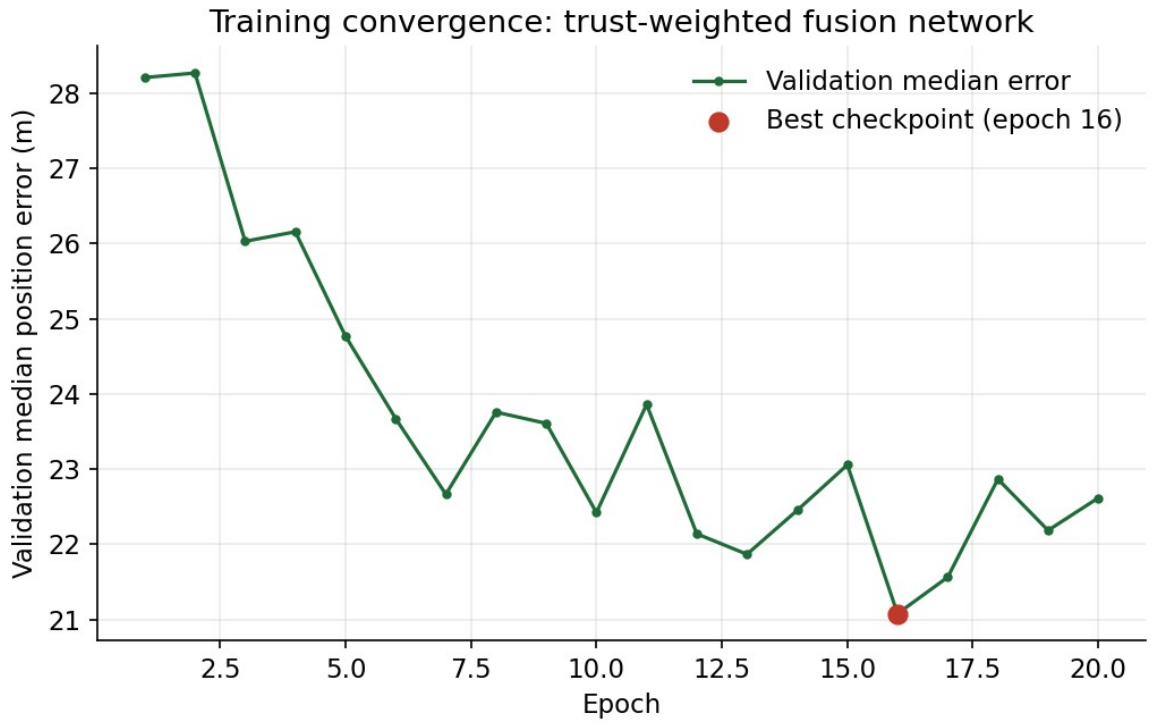

Fig. 5*. Representative training convergence curve, validation median position error by epoch.*

B. Pattern of Improvement

The improvement is not uniform, and its structure is informative. The four largest gains (53–57%) occur in the two scenarios deliberately designed to stress classical geometry distributed and poor_geometry while the smallest gain (7.4%) occurs in formation/SPOOFED, the combination with the highest classical baseline error (70.6 m) and a documented, structural late-run geometric degradation as aircraft converge toward a shared destination under prolonged, undetected spoofing. This pattern demonstrates larger relative improvement precisely where the fixed inverse-covariance weighting scheme is weakest and is consistent with the learned network exploiting context-dependent trust allocation rather than uniformly reweighting all inputs.

C. EKF Baseline and Integrity-Gating Ablation

The plain EKF baseline is competitive with or better than classical WLS under jamming in the richest-geometry scenario (large_fleet), where temporal smoothing has genuine noise to average out, but it is substantially worse than both classical WLS and the learned network under spoofing and most dramatically in formation/SPOOFED, where the EKF's median error (230.8 m) is more than three times the classical baseline's and nearly four times the trained network's. This is attributed to the absence of any integrity-monitoring mechanism in EKF. Its sequential temporal memory allows a slow, undetected spoofing bias to accumulate progressively, whereas the memoryless classical baseline effectively "resets" any accumulated bias at every timestep.

A natural hypothesis was that adding integrity monitoring to the EKF using chi-square innovation gating with persistence-based isolation and reinstatement, as in the framework's own integrity monitor would close this gap. This hypothesis was tested directly and rejected by the ablation results: across both a hard-exclusion variant and a soft

down-weighting variant evaluated at several inflation factors, the gated EKF performed consistently worse than the plain, ungated EKF in the formation/SPOOFED case. Thus, adding a conventional innovation-based integrity monitor did not merely fail to close the gap with the learned network; under this adversarial geometry, it degraded the EKF's performance relative to retaining all measurements. The result indicates that the weakness is not simply the absence of a gating mechanism, but the incompatibility between naive innovation gating and the correlated, partially state-derived measurements present in this cooperative architecture. The most plausible explanation is that naive innovation gating compares each measurement with the filter’s own internal state. However, several cooperative sources are partly derived from the platform’s INS estimate. If a slow spoof has already biased the INS, it can bias these supposedly independent sources in the same direction. The resulting correlation undermines a consistency test that assumes independent measurements. This supports the framework’s architectural choice to detect spoofing by cross-checking against an independently and freshly computed cooperative solution, rather than against the platform’s own state estimate.

D. Three-Phase Denial and Recovery Timeline

The aggregate statistics of Sections VI.A-VI.C establish that the trained network reduces median position error across a large validation set, but do not directly show what an individual denial event looks like over time .What the operator-relevant question actually is that ,how large does the error get, how quickly is it bounded, and how cleanly does the system recover once GPS is restored. These are the logical queries from operations perspectives. To answer these questions directly, three representative scenarios were re-run with a finite denial window (GPS lost at t = 400 s, restored at t = 1000 s, with full run duration being 1800 s) instead of denial persisting to the end of the run as in the training scenarios. This was simulated so that all three phases -- normal operation, active denial, and recovery -- are captured in a single timeline. Because no cooperative-source measurement in the logged data is independent of the platform's own onboard navigation

state absent a cooperative service, the “without cooperative PNT service” baseline shown is a reconstruction rather than a directly logged quantity. It isolates the INS free-drift accumulated specifically during the denial window (relative to the drift-only trajectory's own value at denial onset) and applies it as an offset from the true position at denial onset, with ordinary GPS-quality tracking assumed immediately before and after the window which is consistent with a real GPS receiver reacquiring normally once the signal returns, independent of whether cooperative correction was available during the outage. Each figure's lower figure additionally shows the estimated network clock bias over the same timeline, demonstrating that time synchronization remains bounded independent of the position-domain denial event.

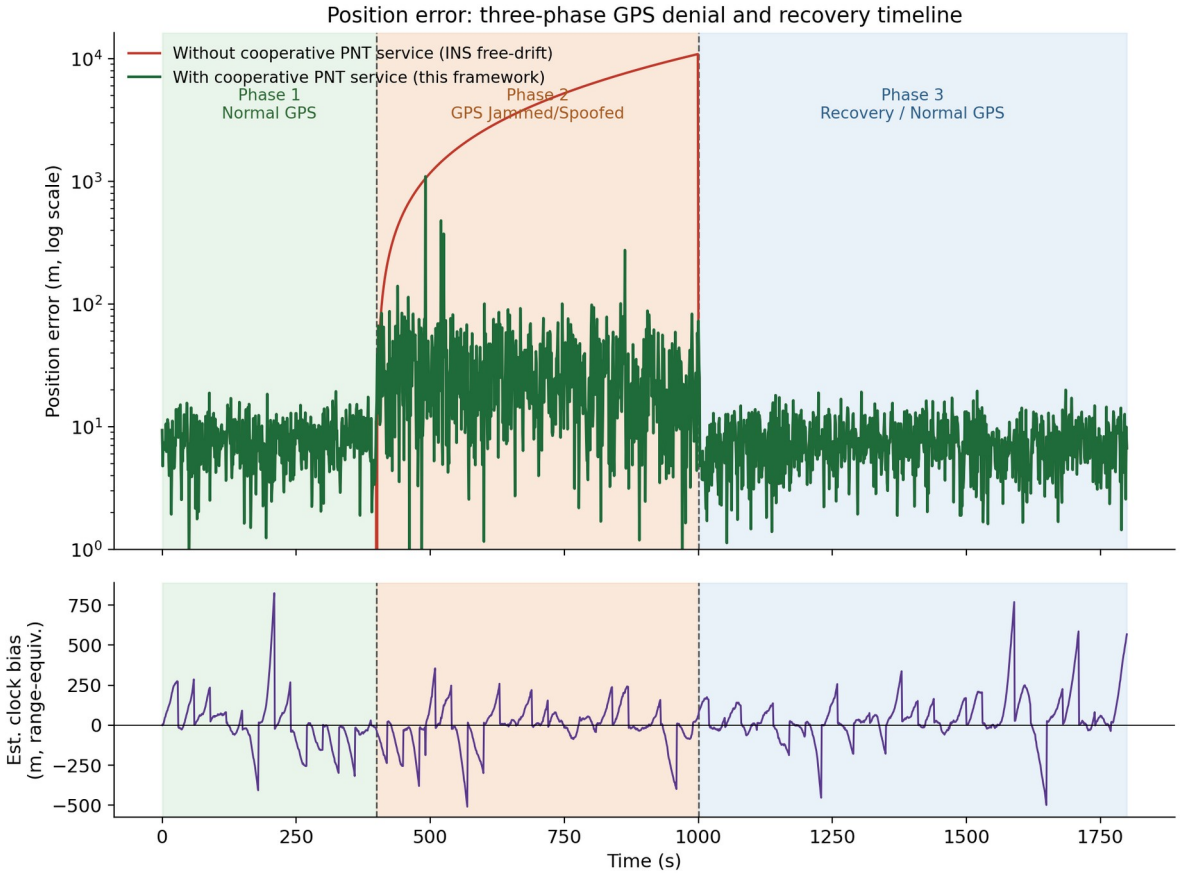

Fig. 6. *Three-phase timeline, formation/JAMMED. Position error (top) stays bounded to tens of metres throughout active denial while the reconstructed without-service baseline grows past 10 km; both curves converge to normal GPS-quality tracking within seconds of*

signal restoration. Estimated clock bias (bottom) remains bounded throughout, unaffected by the position-domain denial event.

Figure 6 shows the formation scenario under GPS jamming conditions, the baseline case against which the harder scenarios are compared throughout this paper. Position error with the cooperative service active remains in the 10-90 m range for the full 600-second denial window, an increase of roughly one order of magnitude over the sub-10 m error typical of normal GPS tracking, but never approaches the kilometre scale. The reconstructed without-service curve, by contrast, grows monotonically and without bound for as long as the denial persists, reaching approximately 11 km by the time GPS is restored at t = 1000 s clearly demonstrating concretely what cooperative fusion is protecting against. Recovery is visually immediate: error returns to the pre-denial baseline within one to two update cycles of GPS reacquisition, consistent with the bump-less transition requirement of automatic recovery function.

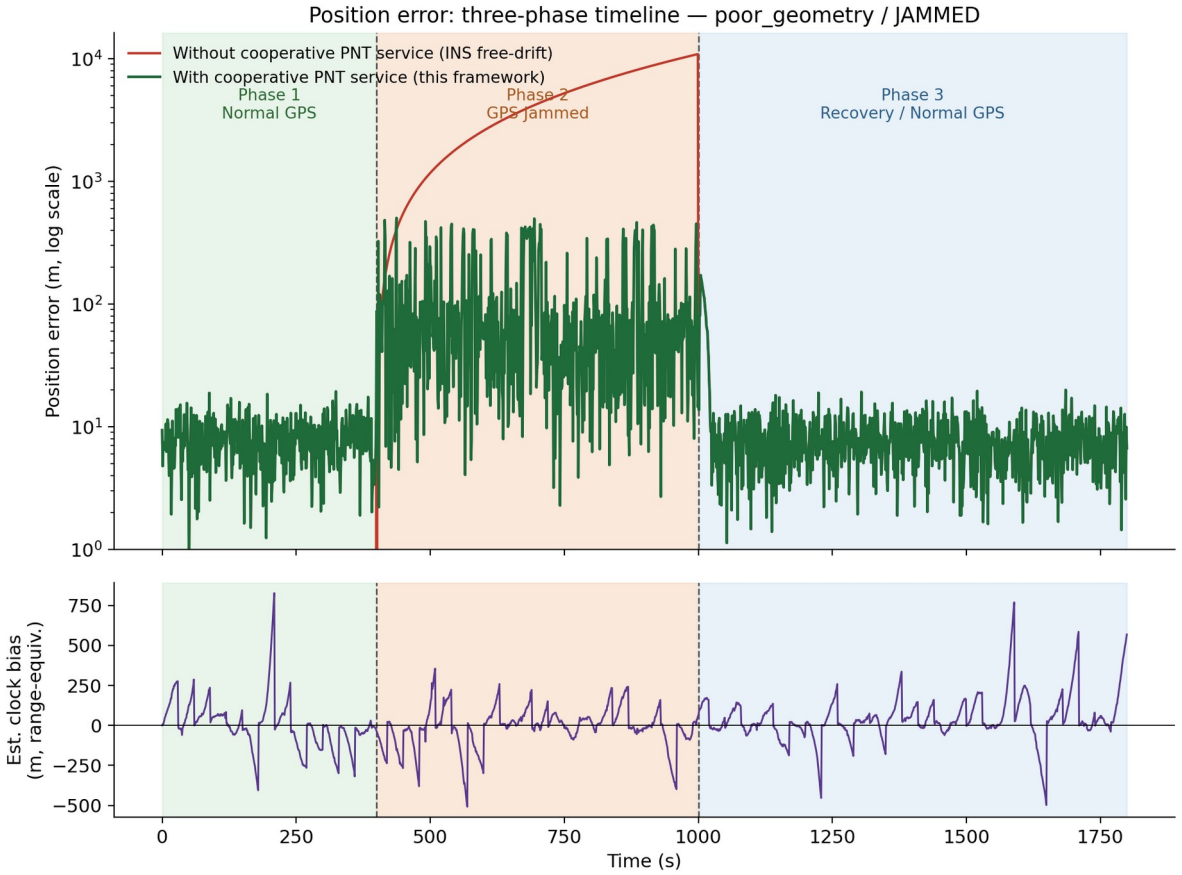

Fig. 7. *Three-phase timeline, poor_geometry/JAMMED -- the deliberately clustered, near-collinear anchor geometry that produced the largest classical-to-learned improvement in Table 2. Position error with the cooperative service remains bounded but visibly noisier and higher in magnitude than the formation case, reaching several hundred metres at its peak, consistent with the higher PDOP this geometry was designed to stress.*

Figure 7 shows the deliberately hard-geometry poor_geometry scenario under jamming, presented as a direct visual contrast to Figure 6. Two differences from the formation case are evident. First, the with-service error during denial is visibly greater and noisier, peaking around near 500 m rather than 90 m directly reflecting the weaker, more clustered anchor geometry this scenario was constructed to represent (mean PDOP 10.4 versus 3.9 for formation, per the scenario design). Secondly, the recovery transition at t = 1000 s is visibly less sharp than in Figure 6, decaying over several update cycles rather than essentially instantaneously, consistent with poorer geometry slowing convergence even after GPS is restored. Despite both effects, the cooperative service still holds error three orders of magnitude below the without-service reconstruction throughout the denial window, and this is precisely the scenario in which the learned network shows its largest measured improvement over classical WLS (Table 2, +57.1% under jamming) . The timeline view and the aggregate statistic further corroborate the findings and tell the same story from two different angles.

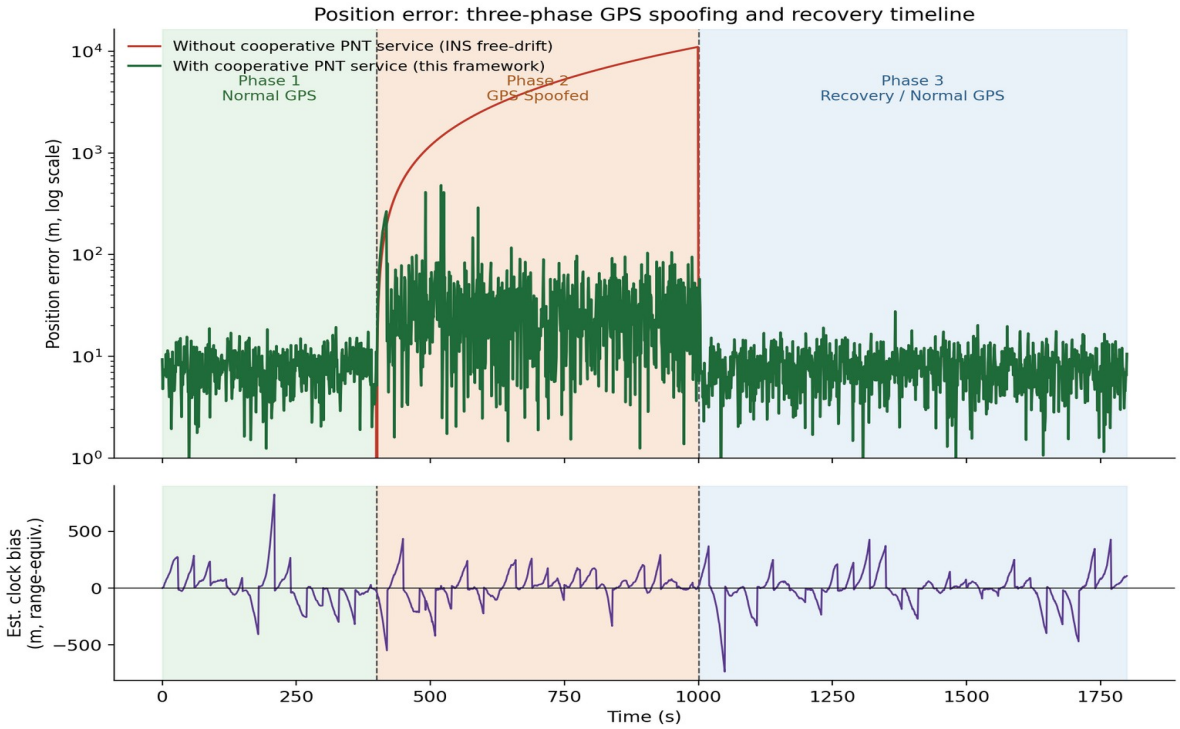

Fig. 8. *Three-phase timeline, formation/SPOOFED -- the single weakest combination identified in Table 2 and Section VI.B, shown here for direct visual grounding of that finding.*

Figure 8 revisits the formation scenario under spoofing rather than jamming conditions, the combination identified in Section VI.B as showing the smallest improvement over classical WLS (+7.4%) and the only combination in which one of five independent training runs failed to improve on the classical baseline at all. The timeline clarifies what that statistic means operationally.The with-service error during denial is visibly higher and less stable than the jammed case in Figure 6, with excursions approaching 500 m rather than remaining under 100 m, consistent with the slow-ramp spoofing bias interacting with the late-run geometric degradation documented elsewhere in this scenario's design. Further more, the system still bounds error more than an order of magnitude below the without-service reconstruction throughout the denial window, and recovery upon GPS restoration remains clean. This figure is included specifically because Table 2 reports only a single summary statistic for this combination; the full timeline shows that the statistic reflects a

genuinely harder, noisier trajectory rather than a training artifact, matching the structural explanation given in Section VI.B.

VII. DISCUSSION

The important experiential finding that a learned, attention-based trust-weighting mechanism outperforms a fixed inverse-covariance scheme by a margin that scales with geometric difficulty has a reasonable mechanistic explanation. Classical inverse-covariance weighting is optimal only in cases wherein each source's reported covariance is itself perfectly calibrated. In practice, a source's reported uncertainty is an approximation (for example, the RSSI-ranging model's realistic but coarse ≈47% relative uncertainty, or a triangulation solution's covariance under near-degenerate bearing geometry), and the degree of miscalibration is plausibly context-dependent in ways a fixed weighting rule cannot capture , however a trained, context-aware attention mechanism can learn to compensate for. This explanation is consistent with, though not definitively proven by, the observed pattern of improvement, and is a natural target for future interpretability analysis of the learned attention weights themselves.

Further analysis was carried out to understand what specifically bounds accuracy in the hardest scenarios, since Table 2's poor_geometry and formation/SPOOFED results remain well above the 10-25 m accuracy typical of nominal INS/GPS operation even with the learned network active. Direct comparison of sensor models shows that SDR two-way ranging is extremely precise across the operational range envelope (sub-metre at 5 km, ~6 m at 200 km), while bearing-based and RSSI-based fixes are inherently coarse at range (a 1.5° bearing uncertainty alone corresponds to a 1.3 km cross-range error at 50 km) . Therefore , a mathematically ideal inverse-covariance scheme would assign bearing and RSSI sources negligible weight whenever any ranging-based multilateration is available. Neither anchor count nor instantaneous PDOP, however, was found to correlate with the largest observed errors in poor_geometry (mean anchors used was statistically

indistinguishable between low-error and high-error steps, and mean PDOP was, counter-intuitively, lower during the worst-error steps than during typical ones). The actual mechanism was identified by re-solving one such worst-case step from the TRUE position as the Gauss-Newton initial guess, eliminating any possibility of poor convergence or a bad initial estimate. The solver still converged to a position 348.5 m from truth. This demonstrates that the observed error is not a defect in the fusion process or a failure of convergence, but rather an inherent characteristic of the data under near-singular geometric conditions. While Position Dilution of Precision (PDOP) describes the average amplification of measurement errors, poorly conditioned geometries can exhibit a heavy-tailed error response. Under such conditions, even small and individually well-calibrated errors from multiple anchors may, by chance, become aligned with the geometry's high-gain direction. This can result in a large position error that is nevertheless mathematically consistent with, and indeed the correct least-squares solution for, the measurements available at that particular time step. This finding was further examined using Iteratively Reweighted Least Squares (IRLS) with Huber weighting. The method reduces the influence of an anchor when its residual is unusually large compared with the expected measurement noise. When applied to the complete poor-geometry/JAMMED denial period, IRLS reduced the maximum position error by about 30%, from 2406.8 m to 1691.4 m. However, the median error remained unchanged at 59.9 m, and the 95th-percentile error changed only slightly, from 453.0 m to 447.4 m.

These results are significant because IRLS is effective mainly when one or more measurements behave as clear outliers. In the observed cases, however, the large errors were primarily caused by geometric amplification of small, normally distributed errors across multiple anchors. Since no individual anchor could be clearly identified as an outlier, there was little for IRLS to down-weight. It therefore improved only the extreme tail of the error distribution, where the effect became large enough for individual measurements to appear anomalous.

Overall, these results indicate that anchor geometry, rather than the choice of fusion algorithm (classical or learned)is the main factor limiting accuracy in the most challenging scenarios. Improving the anchor geometry is therefore the most direct way to reduce these errors and approach nominal INS/GPS-level accuracy. Temporal smoothing, discussed in Section VIII, may provide an additional mitigation method because these large errors are transient, although it was not implemented in the present study. This was the deliberate decision as it is felt that it can also open up a spoofing window

The EKF findings carry a distinct implication i.e the common assumption that a sequential filter is a strictly stronger baseline than a memoryless per-step solution does not hold once an adversarial (rather than merely noisy) denial condition is considered, unless the filter's integrity monitoring is designed with the same care given to the rest of the estimation architecture. This is a cautionary result for any GPS-denied cooperative navigation design that adopts a Kalman-filter-family estimator without an explicitly independent integrity reference.

The single honest exception to an otherwise uniformly positive result is formation/SPOOFED, where one of five independent training runs showed no improvement over classical WLS . This is attributable to a structural characteristic of that specific scenario (late-run geometric degradation as aircraft converge toward a shared destination under prolonged undetected spoofing, affecting roughly 1.6% of denial-step records across the full dataset) rather than to general training instability, since the remaining nine combinations were robustly positive across all five runs. We are honestly reporting this exception rather than omittting it, consistent with the paper's broader emphasis on statistically grounded, honestly scoped claims.

VIII. LIMITATIONS AND FUTURE WORK

(1) Validation is simulation-only. No hardware-in-the-loop or flight-test validation has been performed. This requirements specification would be considered as a necessary follow-on work for future.

(2) “Secure” cooperative broadcast is modeled as a received, authenticated output rather than an implemented cryptographic and RF protocol stack; IFF identification uses a numeric code exchange rather than a Mode-5 cryptographic protocol. Both are explicitly out of scope of this work and should not be read as security validation.

(3) Satellite visibility / constellation-geometry monitoring has not been modeled. GPS availability is treated as a scenario-level input rather than derived from a simulated constellation.

(4) The known-emitter database for RWR/ESM geolocation is a static, scenario-configured list rather than a dynamically managed, versioned asset.

(5) The trust-weighted fusion network is currently memoryless, freshly computing trust weights from first principles at every timestep. A natural extension inspired by the EKF's large_fleet result, where temporal smoothing helped under non-adversarial conditions is to add a persistent, learned temporal state, extending the GRU-based persistent embedding used for temporal track association in our prior work [7], so the network can learn when to exploit history and when to discount it, potentially obtaining the EKF's smoothing benefit without its demonstrated spoofing vulnerability. This would be taken up subsequently.

(6) A Federated Kalman Filter and a factor-graph-based smoother were not implemented as additional baselines in this study and remain natural candidates for future comparison [10–12].

(7) Statistical validation, while based on five independent training runs, could be further strengthened with additional independent runs and formal significance testing beyond the mean ± standard deviation reporting used here.

IX. CONCLUSION

This paper presented a full-stack, simulation-validated framework for Collaborative PNT over a tactical SDR network, and a learned, attention-based trust-weighted fusion network that extends a prior radar/IFF track-fusion architecture to the variable-cardinality navigation-source fusion problem. Across a statistically grounded, five-independent-run evaluation spanning five scenario geometries and both jamming and spoofing denial, the learned network reduced median position error by 7–57% relative to a classical weighted-least-squares baseline, with the largest gains in the geometrically hardest scenarios. A literature-standard EKF baseline was found to be substantially more vulnerable to spoofing than either the classical or learned alternatives, and a follow-up ablation showed that naive integrity gating does not resolve this vulnerability without an independently-referenced consistency check reinforcing the architectural rationale for the framework's own integrity-monitoring design. These results substantiate the claims that learned, context-adaptive trust weighting is a genuine and measurable improvement over fixed-weighting classical fusion for cooperative GPS-denied navigation, while honestly identifying the scope, limitations, and open questions that remain before its operational deployment.

X. REFERENCES

[1] Jarraya, I., Al-Batati, A., Kadri, M. B., Abdelkader, M., Ammar, A., Boulila, W., and Koubâa, A., “GNSS-Denied Unmanned Aerial Vehicle Navigation: Analyzing Computational Complexity, Sensor Fusion, and Localization Methodologies,” Satellite Navigation, Vol. 6, No. 1, 2025. https://doi.org/10.1186/s43020-025-00162-z

[2] Gabela, J., Retscher, G., Goel, S., Perakis, H., Masiero, A., Toth, C., Gikas, V., Kealy, A., Koppányi, Z., Błaszczak-Bąk, W., Li, Y., and Grejner-Brzezinska, D. A., “Experimental Evaluation of a UWB-Based Cooperative Positioning System for Pedestrians in GNSS-Denied Environment,” Sensors, Vol. 19, No. 23, 2019, Art. 5274. https://doi.org/10.3390/s19235274

[3] Jia, C., Ren, P., and Deng, T., “A Novel Ranging and IMU-Based Method for Relative Positioning of Two-MAV Formation in GNSS-Denied Environments,” Sensors, Vol. 23, No. 9, 2023, Art. 4366. https://doi.org/10.3390/s23094366

[4] Ellingson, G., Brink, K., and McLain, T., “Relative Navigation of Fixed-Wing Aircraft in GPS-Denied Environments,” NAVIGATION: Journal of the Institute of Navigation, Vol. 67, No. 2, 2020, pp. 255–273. https://doi.org/10.1002/navi.364

[5] Ellingson, G., Brink, K., and McLain, T., “Cooperative Relative Navigation of Multiple Aircraft in Global Positioning System-Denied/Degraded Environments,” Journal of Aerospace Information Systems, 2021. https://doi.org/10.2514/1.I010802

[6] Shen, J., Wang, S., Zhai, Y., and Zhan, X., “Cooperative Relative Navigation for Multi-UAV Systems by Exploiting GNSS and Peer-to-Peer Ranging Measurements,” IET Radar, Sonar & Navigation, Vol. 15, No. 1, 2021, pp. 21–36. https://doi.org/10.1049/rsn2.12023

[7] Sharma, G., Kumar, C. R. S., and Rajesh, R., “Multi-Sensor Data Fusion for Aerial Target Tracking Using Transformer-Based Association Attention and Self-Supervised Fine-Tuning,” Proceedings of the IEEE MAPCON Conference, Nagpur, India, Dec. 2026 [under review].

[8] Sharma, G., Kumar, C. R. S., and Rajesh, R., “Transformer-Based Multi-Sensor Data Fusion for Airborne Surveillance Systems Using Adaptive Covariance Intersection,” Journal of Aerospace Information Systems [under review].

[9] Julier, S. J., and Uhlmann, J. K., “General Decentralized Data Fusion with Covariance Intersection,” Handbook of Multisensor Data Fusion: Theory and Practice, CRC Press, 2009, pp. 319–344.

[10] Yan, Y., and Yang, J., “An Improved Variational Bayesian-Based Adaptive Federated Kalman Filter for Multi-Sensor Integrated Navigation Systems,” Sensors, Vol. 25, No. 23, 2025. https://doi.org/10.3390/s25237173

[11] Zhao, Q., Lü, S., Lou, J., and Zhang, R., “FE-GUT: Factor Graph Optimization Hybrid with Extended Kalman Filter for Tightly Coupled GNSS/UWB Integration,” arXiv:2407.06915, 2024.

[12] Xu, Z., Zhu, F., and Zhang, X., “State Estimation in Multi-Sensor Fusion Navigation: Equivalence Analysis on Filtering and Optimization,” International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XLVIII-1/W2-2023, 2023, pp. 1161–1168. https://doi.org/10.5194/isprs-archives-XLVIII-1-W2-2023-1161-2023

[13] Abdolkarimi, E. S., and Rafatnia, S., “Fuzzy Fixed-Time Multi-Source Fusion with Attention-Based Neural Networks for Reliable Navigation under Satellite Signal Loss,” Engineering Applications of Artificial Intelligence, Vol. 163, 2025, Art. 112806. https://doi.org/10.1016/j.engappai.2025.112806

[14] Yang, B., Wang, X., Huang, F., Cao, X., and Zhang, Z., “R-AFNIO: Redundant IMU Fusion with Attention Mechanism for Neural Inertial Odometry,” Expert Systems with Applications, Vol. 265, 2024, Art. 125894. https://doi.org/10.1016/j.eswa.2024.125894

[15] Tanil, C., Khanafseh, S., Joerger, M., and Pervan, B., “Spoofing Detection Using RAIM with INS Coupling,” Proceedings of the IEEE/ION Position, Location and Navigation Symposium (PLANS), IEEE, 2014.

[16] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I., “Attention Is All You Need,” Advances in Neural Information Processing Systems 30 (NeurIPS 2017), 2017, pp. 5998–6008.